arXiv:1603.05522v1 [stat.AP] 17 Mar 2016

Tracking multiple moving objects in images using Markov Chain Monte Carlo

Lan Jiang & Sumeetpal S. Singh *†

March 18, 2016

Abstract

A new Bayesian state and parameter learning algorithm for multiple target tracking (MTT) models with image observations is proposed. Specifically, a Markov chain Monte Carlo algorithm is designed to sample from the posterior distribution of the unknown number of targets, their birth and death times, states and model parameters, which constitutes the complete solution to the tracking problem. The conventional approach is to pre-process the images to extract point observations and then perform tracking. We model the image generation process directly to avoid potential loss of information when extracting point observations. Numerical examples show that our algorithm has improved tracking performance over commonly used techniques, for both synthetic examples and real florescent microscopy data, especially in the case of dim targets with overlapping illuminated regions.

1 Introduction

The multiple target tracking (MTT) problem is to infer the states or tracks of multiple moving objects from noisy measurements. The problem is difficult since the number of targets is unknown and changes over time as it is a birth-death process. Other compounding factors include the non-linearity of both the target's motion and observation models. In many applications such as radar/sonar tracking [1] and Fluorescence Microscopy [2], the measurements (or observations) are images. (For example, a pixel's illumination intensity is a measure of nearby targets energy and background noise.) These images are usually pre-processed prior to actual tracking to extract point measurements where each point is a spatial coordinate, which are then assumed to be either noisy measurements of the target state or spuriously generated. The latter is an artefact of the method that extracts point measurements. Converting images to point measurements is advantageous because it yields a simpler observation model and also simplifies the design of tracking algorithms [1, 3], e.g. [2] connects the point measurements using a nearest neighbour method to form target trajectories. However, the pre-processing step can introduce information loss in the low signal-to-noise (SNR) regime, which can be near complete as the targets become more closely spaced and background noise intensifies. In low SNR it can be difficult to isolate bright regions in the image, whose centres would be the candidate point measurements, and then attribute them to distinct targets. Thus, MTT algorithms that work with the images directly can be preferable and there is a sizeable literature on it. A selection

*The authors are at the Department of Engineering, University of Cambridge, United Kingdom.

†This work was supported by the Engineering and Physical Sciences Research Council [grant numbers EP/G037590/1, EP/K020153/1.]

of works is [4, 5, 6, 7, 8, 9, 10] and they differ in how tracking is achieved (Bayesian, maximum likelihood or otherwise) and the specific assumptions imposed on the image model.

Given images recorded over a length of time, say from time 1 to n, our aim is to jointly infer the target tracks and the MTT model parameters. We adopt a Bayesian approach and one of our main contributions is the design of a new Markov chain Monte Carlo (MCMC) algorithm for an image measurement model that can jointly track and calibrate. The MCMC algorithm is a trans-dimensional sampler that combines Particle Markov Chain Monte Carlo (PMCMC) steps [11] to sample from the *exact* MTT posterior distribution for entire target tracks and model parameters. This is in contrast to numerous techniques that use specific variational approximations, e.g. spatial Poisson, of the MTT posterior to simplify inference [3]. Entire tracks, as opposed to point estimates of target locations at each time [9], are needed as these are then used to infer the aggregate diffusion characteristics of the molecules [2]. Model calibration is needed because in real fluorescence microscopy data, the molecules to be tracked bleach over time and the noise characteristics of the acquired images also drift. These changes can be captured by time varying image model parameters, which are indeed unknown to the analyst, as are other parameters such as those describing the molecule motion model. To the best of our knowledge, our trans-dimensional MCMC tracker for image observations addresses this practical tracking problem in greater generality without major limiting assumptions like well separated molecules with non-overlapping illumination regions [9] or aforementioned principled simplifications of the MTT posterior [10]. In numerical examples we demonstrate the superior performance of our method over [9] for closely spaced targets. We also apply it to real florescence microscopy data and show it outperforms a method currently used by biologists [2] which pre-processes the images to extract point observations. Such comparisons, which are absent in the literature, highlight the gain in performance by targeting the exact MTT posterior and avoiding simplifications like disallowing overlaps.

We do not advocate that our MCMC technique should replace techniques that extract point observations, those that use variational approximations to simplify the posterior [3, 9, 10], or those that do not extract point observations but are optimised for non-overlapping targets [9]. Our MCMC technique should be viewed as a compliment to these other techniques. It could be applied in an online tracking scenario by processing a window of data at a time or as a post-processing tool to refine the trajectories identified by any other online algorithm [12]. This is similar to the role MCMC plays in the related field of Particle Filtering, which is an online estimation method, where MCMC is used to refine the online estimates [13].

There is a growing literature on using MCMC for tracking as it is recognised that sampling the true MTT posterior, although challenging, is feasible in offline applications and can serve as a track refinement tool in the online setting [14, 15]. There exists several MCMC based MTT algorithms for point observation models. [14, 16] assume the underlying state-space and observation model is linear and Gaussian, [15, 17] consider the non-linear and non-Gaussian setting while [16, 17, 18] simultaneously estimate the model parameters. (Although some of the above works incorporate parameter estimation, it is a topic in MTT that has only recently gained attention, see [19, 20].) Tracking using images are also known as track-before-detect (TBD) techniques. [7, 9, 10] use specific but different Poisson approximations for the MTT posterior (assuming known model parameters) which is then approximated using a Particle Filter.

The remainder of the paper is organised as follows. Section 2 describes the MTT model and presents the framework for joint state and parameter learning algorithm. In Section 3, we present details of our novel MCMC kernel for detecting and maintaining tracks, which constitutes the core part of our tracking algorithm. (More detailed derivations are given in the Appendix.) Section 4 presents numerical results for both synthetic and real florescent microscopy data.

2 Multiple target tracking model

2.1 The single target model

We commence with a description of the image based tracking problem assuming a single target and then enlarge the model for the multi-target case. Let the Markov process $\{X_t\}_{t\geq 1}$ represent the state values of a single evolving target. In this work it is assumed that $X_t = (X_t(1), \ldots, X_t(5)) \in \mathbb{R}^5$ where $X_t(i)$ denotes its ith component. $X_t(1)$ is the target's *illumination intensity* or amplitude (to be discussed in detail next), $(X_t(2), X_t(3))$ is the spatial coordinate of the target and $(X_t(4), X_t(5))$ are the corresponding spatial velocities. Frequent reference will be made to the intensity (amplitude), spatial coordinate and spatial velocity components of a target state X_t. As such we will denote these components by $A_t = X_t(1)$, $S_t = (S_t(1), S_t(2)) = (X_t(2), X_t(3))$ and $V_t = (V_t(1), V_t(2)) = (X_t(4), X_t(5))$. X_t is a time-homogeneous Markov process,

$$X_1 \sim \mu_\psi(\cdot), \quad X_t | X_{1:t-1} = x_{1:t-1} \sim f_\psi(\cdot|x_{t-1}) \tag{1}$$

where μ_ψ and f_ψ are, respectively, the initial and state transition probability density function (pdf), both parametrised by the common real valued vector $\psi \in \Psi \subset \mathbb{R}^{d_\psi}$. (As a rule, a random variable (r.v.) is denoted by a capital letter and its realisation by small case.) For example, for linear and Gaussian state dynamics, $\mu_\psi(x) = \mathcal{N}(x; \mu_b, \Sigma_b)$, $f_\psi(x'|x) = \mathcal{N}(x'; Fx, W)$, where $\mathcal{N}(\cdot; m, \Sigma)$ denotes the Gaussian pdf with mean m and covariance matrix Σ. Thus $\psi = (\mu_b, \Sigma_b, F, W)$.

For the measurements, a two dimensional image measurement model is assumed with m pixels in total. Let

$$Y_t = (Y_{t,1}, \ldots, Y_{t,m}),$$

denote the observed image at time t where $Y_{t,i}$ is the value (illumination intensity) of pixel i. $Y_{t,i}$ is defined as

$$Y_{t,i} = h_i(X_t) + E_{t,i}, \tag{2}$$

where $E_{t,i}$ is the observation noise of pixel i at time t and $h_i(X_t)$ is the illumination of pixel i by a single target with state X_t. As in [9], for $x = (a, s, v) \in \mathbb{R} \times \mathbb{R}^2 \times \mathbb{R}^2$ where a is the intensity, $s = (s(1), s(2))$ the spatial coordinate and $v = (v(1), v(2))$ the spatial velocity, $h_i(x)$ is the point spread function

$$\begin{aligned} h_i(x) &= \mathbb{I}[i \in L(s)] \, \frac{a\Delta_1\Delta_2}{2\pi\sigma_h^2} \\ &\quad \times \exp\{-\frac{(\Delta_1 r - s(1))^2 + (\Delta_2 c - s(2))^2}{2\sigma_h^2}\} \\ &=: \; a \, \bar{h}_i(s) \end{aligned} \tag{3}$$

where (r, c) denotes the row and column number of pixel i, Δ_1 and Δ_2 are constants that map pixel indices to spatial coordinates and σ_h is the blurring parameter. It is assumed that the spatial coordinate of the pixel with index corresponding to row and column number $(0, 0)$ is the origin of $\mathbb{R}^2$. As in [9], we also assume for each state value $x = (a, s, v)$ there is a square truncation region $L(s)$ where $h_i(x) = 0$ if $i \notin L(s)$. Specifically, $L(s)$ is the set of $l \times l$ pixels, l an odd integer, whose centre pixel has spatial coordinate closest to s. Henceforth we assume $\Delta_1 = \Delta_2 = \Delta$.

For later use, the function $\bar{h}_i(s)$ in (3) has been implicitly defined. In addition, extend the domain of the truncation region L and point spread function $\bar{h}_i$ to included pixel indices

$j \in \{1, \ldots, m\}$. That is, let (r', c') be the row and column number of pixel j and define

$$L(j) = L(s), \quad \bar{h}_i(j) = \bar{h}_i(s) \quad \text{where } s = (r'\Delta, c'\Delta). \tag{4}$$

Equivalently $L(j)$ is the square of $l \times l$ pixels centered at pixel j.

The pixel noise is assumed to be Gaussian with mean value b_t, representing the background intensity, and variance $\sigma_{r,t}^2$, both time varying but common across pixels, i.e.[1]

$$E_{t,i} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\cdot|b_t, \sigma_{r,t}^2), \quad i = 1, \ldots, m, \ t = 1, \ldots, n.$$

Thus, the conditional pdf of the observed image at time t due to a single target with state X_t is

$$g_t(y_t|x_t) = \prod_{i=1}^{m} \mathcal{N}(y_{t,i}; h_i(x_t) + b_t, \sigma_{r,t}^2).$$

where subscript t of g_t indicates the observation model is time-inhomogeneous. Given n images, all the model parameters $(\psi, b_1, \sigma_{r,1}, \ldots, b_n, \sigma_{r,n})$ described in this section will be estimated.

2.2 The model for multiple targets

In this section we partially adopt the formulation in [17] for the MTT model. (Note though that the observation model in [17] is for point-observations and not for images as in our case.) In an MTT model, the MTT state at time t is the concatenation of all individual target states at t:

$$\mathbf{X}_t = \left(X_{t,1}, X_{t,2}, \ldots, X_{t,K_t^x}\right)$$

where each sub-vector $X_{t,i}$ is the state (as in (1)) of an individual target. The number of targets K_t^x under surveillance changes over time due to the death of existing targets and the birth of new targets. Independently of the other targets, a target survives to the next time with survival probability p_s and its state evolves according to the transition density f_ψ, otherwise it 'dies'. In addition to the surviving targets, new targets are 'born' from a Poisson process with intensity λ_b and each of their states is initialised by sampling from the initial density μ_ψ. The states of the new born targets and surviving targets from time t make up $\mathbf{X}_{t+1}$. We assume that at time $t = 1$ there are only new born targets, i.e. no surviving targets from the past.

To describe the evolution of $\mathbf{X}_t$ due to survivals and births, a series of random variables are now defined here. Let K_t^x and K_t^b denote the number of targets and new births at time t respectively. We start with $K_1^x = K_1^b$. For $t > 1$ and $i = 1, \ldots, K_{t-1}^x$, let

$$C_t(i) = \begin{cases} 1 & i\text{th target at time } t-1 \text{ survives to time } t \\ 0 & i\text{th target at time } t-1 \text{ does not survive to } t \end{cases}.$$

C_t is the $K_{t-1}^x \times 1$ binary vector where 1's indicate survivals and 0's indicate deaths of targets from time $t-1$. Let K_t^s denote the number of surviving targets at time t, thus

$$K_t^s = \sum_{i=1}^{K_{t-1}^x} C_t(i).$$

[1] $\sigma_{r,t}$ is the symbol for the observation noise and subscript r is not to be confused with row number mentioned before.

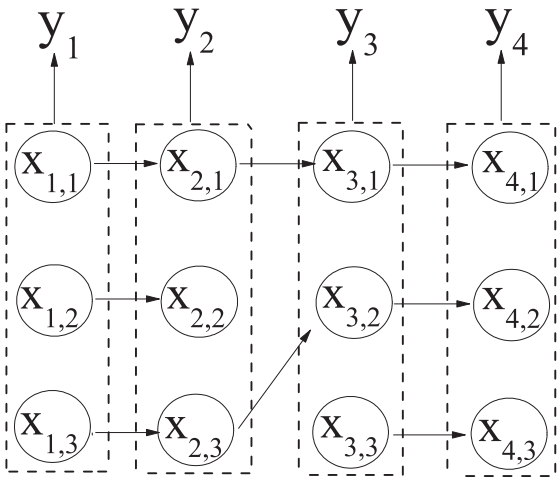

Figure 1: A realisation from the MTT model: states of a target are connected with arrows and all targets at time t contribute to image y_t.
MTT random variables:
Time $t = 1$: No prior targets ($C_1 = (), K_1^s = 0, I_1 = ()$), three targets are born ($K_1^x = K_1^b = 3$) with states $X_{1,1}, X_{1,2}, X_{1,3}$;
Time $t = 2$: All targets $X_{1,1}, X_{1,2}, X_{1,3}$ survive to $X_{2,1}, X_{2,2}, X_{2,3}$. Thus $C_2 = (1, 1, 1)$, $K_2^s = 3$, $I_2 = (1, 2, 3)$. No new born targets, $K_2^b = 0$, $K_2^x = K_2^s + K_2^b = 3$.
Time $t = 3$: Targets $X_{2,1}$ and $X_{2,3}$ survive to $X_{3,1}$ and $X_{3,2}$ respectively while $X_{2,2}$ dies, thus $C_3 = (1, 0, 1)$, $K_3^s = 2$, $I_3 = (1, 3)$. One new born target, $K_3^b = 1$, denoted $X_{3,3}$. $K_3^x = K_3^s + K_3^b = 3$.
Time $t = 4$: All targets survive, no new born, same as time $t = 2$.
MTT variables of the equivalent description of Sec. 2.4:
$t_b^1 = 1$, $\mathbf{X}^1 = (X_{1,1}, X_{2,1}, X_{3,1}, X_{4,1})$; $t_b^2 = 1$, $\mathbf{X}^2 = (X_{1,2}, X_{2,2})$; $t_b^3 = 1$, $\mathbf{X}^3 = (X_{1,3}, X_{2,3}, X_{3,2}, X_{4,2})$; $t_b^4 = 3$, $\mathbf{X}^4 = (X_{3,3}, X_{4,3})$.

The K_t^s surviving targets from time $t-1$ evolve to become the *first* K_t^s targets in $\mathbf{X}_t$. Specifically, define the $K_t^s \times 1$ ancestor vector I_t, $t > 1$, as

$$I_t(i) = \min\{k : \sum_{j=1}^{k} C_t(j) = i\}, \quad i = 1, \ldots, K_t^s,$$

so that $X_{t-1,I_t(i)}$ evolves to $X_{t,i}$ for $i = 1, \ldots, K_t^s$. In addition to the surviving targets $(X_{t,1}, \ldots, X_{t,K_t^s})$, we have K_t^b newly born targets denoted by $X_{t,K_t^s+1}, \ldots, X_{t,K_t^x}$. The state $\mathbf{X}_t$ is formed of the new born targets together with the surviving targets, and thus $K_t^x = K_t^b + K_t^s$. An ordering rule is adopted for the new born targets to avoid labelling ambiguity. Specifically, the new born targets at each time t are labelled in ascending order of their first component value. Let $Z_1 = K_1^b$ and

$$Z_t = \left(C_t, K_t^b\right), \quad t > 1,$$

which is the discrete component of the MTT state at time t. Figure 1 illustrates all the MTT random variables.

2.3 The law of MTT model

The image observation $Y_t = (Y_{t,1}, \ldots, Y_{t,m})$ generated by multiple targets at time t is the superposition of the contributions of all targets at time t, the background intensity and noise,

i.e.

$$Y_{t,i} = h_i(\mathbf{X}_t) + E_{t,i}, \quad h_i(\mathbf{X}_t) = \sum_{k=1}^{K_t^x} h_i(X_{t,k}), \tag{5}$$

where $h_i(X_{t,k})$ is the contribution of the k-th target at time t to the illumination of pixel i (see (3)). The MTT observation model is

$$g_t(y_t|\mathbf{x}_t) = \prod_{i=1}^{m} \mathcal{N}(y_{t,i}; h_i(\mathbf{x}_t) + b_t, \sigma_{r,t}^2). \tag{6}$$

Given the vector of the MTT model parameters

$$\theta = (\psi, p_s, \lambda_b, b_1, \sigma_{r,1}^2, \ldots, b_n, \sigma_{r,n}^2) \tag{7}$$

the law of the MTT model can be expressed with the joint density of $(Z_{1:n}, \mathbf{X}_{1:n}, Y_{1:n})$ which is

$$p_\theta(z_{1:n}, \mathbf{x}_{1:n}, y_{1:n}) = p_\theta(y_{1:n}|\mathbf{x}_{1:n}, z_{1:n}) p_\theta(\mathbf{x}_{1:n}|z_{1:n}) p_\theta(z_{1:n})$$

where $a_{i:j}$, $i \leq j$, denotes the sequence $a_i, a_{i+1} \ldots a_j$.

$$p_\theta(y_{1:n}|\mathbf{x}_{1:n}, z_{1:n}) = \prod_{t=1}^{n} g_\theta(y_t|\mathbf{x}_t), \tag{8}$$

$$p_\theta(z_{1:n}) = \mathcal{P}(k_1^b; \lambda_b) \prod_{t=2}^{n} p_s^{k_t^s} (1-p_s)^{k_{t-1}^x - k_t^s} \mathcal{P}(k_t^b; \lambda_b), \tag{9}$$

$$p_\theta(\mathbf{x}_{1:n}|z_{1:n}) = \prod_{t=1}^{n} \Bigg[\prod_{j=1}^{k_t^s} f_\psi(x_{t,j}|x_{t-1,i_t(j)}) \; k_t^b! \mathbb{I}_{\mathcal{O}}(x_{t,k_t^s+1:k_t^x}) \prod_{j=k_t^s+1}^{k_t^x} \mu_\psi(x_{t,j}) \Bigg] \tag{10}$$

where $\mathcal{P}(k;\lambda)$ denotes the probability mass function of the Poisson distribution with mean λ and $g_\theta(y_t|\mathbf{x}_t)$ in (8) is only dependent on components $(b_t, \sigma_{r,t})$ of θ and is precisely g_t of (6). In (10), $\mathbb{I}_{\mathcal{O}}$ is the indicator function of the particular ordering rule $\mathcal{O}$ for the new born targets,

$$\mathbb{I}_{\mathcal{O}}(x_{t,k_t^s+1:k_t^x}) = \begin{cases} 1 & \text{if } x_{t,k_t^s+1}(1) < \cdots < x_{t,k_t^x}(1), \\ 0 & \text{else.} \end{cases}$$

Note that ordering the latent variables is also done in other statistical problems where labelling ambiguity arises through the likelihood function, e.g. Bayesian inference of mixture distributions [21]. For MTT, the ordering will be very useful in Section 2.4 where, thanks to rule $\mathcal{O}$, we are able to deterministically assign a unique label to each target track. Finally, the marginal likelihood of the data $y_{1:n}$ is given by

$$p_\theta(y_{1:n}) = \sum_{z_{1:n}} p_\theta(z_{1:n}) \int p_\theta(y_{1:n}|\mathbf{x}_{1:n}, z_{1:n}) p_\theta(\mathbf{x}_{1:n}|z_{1:n}) \mathrm{d}\mathbf{x}_{1:n}.$$

2.4 An equivalent representation of $(Z_{1:n}, \mathbf{X}_{1:n})$

This section introduces an equivalent parameterization for the MTT problem. Essentially, we define a new set of random variables which are an alternative to those defined in Section 2.3 without any loss of information. The idea here is to introduce notation that explicitly isolates the state trajectories of individual targets, which will be very useful to describe the MCMC proposal distributions in Section 3.

Let $K = \sum_{t=1}^{n} k_t^b$ denote the total number of targets that have appeared from time 1 to n. Each target appearing in this time span can be assigned a distinct label or index $k \in \{1, \ldots, K\}$ with the convention that targets born earlier are given a smaller label than those born at a later time and targets born at the same time are sorted by the ordering rule $\mathcal{O}$.

Consider a target assigned labelled $k \in \{1, \ldots, K\}$, let it's birth time be t_b^k, death time be t_d^k and its life span be $l^k = t_d^k - t_b^k$. (Note $t_d^k - 1$ is the final time of its existence.) The entire continuous state trajectory of this target can be extracted from the MTT state sequence $(\mathbf{X}_{t_b^k}, \ldots, \mathbf{X}_{t_d^k - 1})$ and denote it by

$$\mathbf{X}^k = (X_0^k, \ldots, X_{l_k - 1}^k)$$

where X_{i-1}^k is the i-th state of target k. Note that $\mathbf{X}^k$ is a Markov process with initial and state transition densities μ_ψ and f_ψ respectively. It is straightforward to extract $\{(k, t_b^k, \mathbf{X}^k)\}_{k=1}^K$ from $(Z_{1:n}, \mathbf{X}_{1:n})$ as illustrated in Figure 1. The main point is that we can use one of the two equivalent descriptions for latent variables of the MTT model, i.e.

$$(Z_{1:n}, \mathbf{X}_{1:n}) \Leftrightarrow \{(k, t_b^k, \mathbf{X}^k)\}_{k=1}^K. \tag{11}$$

On the other hand, $(Z_{1:n}, \mathbf{X}_{1:n})$ can be obtained from $\{(k, t_b^k, \mathbf{X}^k)\}_{k=1}^K$ since the underlying transformation is a bijection. (Again see Figure 1 for an example.)

2.5 Bayesian tracking and parameter estimation for MTT

The inference task is to estimate the discrete variables $Z_{1:n}$, target states $\mathbf{X}_{1:n}$ and the MTT parameter θ given the observations $y_{1:n}$. Regarding θ as a random variable taking values in Θ with a prior density $\eta(\theta)$, the goal is to obtain Monte Carlo samples from

$$p(z_{1:n}, \mathbf{x}_{1:n}, \theta | y_{1:n}) \propto \eta(\theta) p_\theta(z_{1:n}, \mathbf{x}_{1:n}, y_{1:n}). \tag{12}$$

We achieve this by iteratively performing the MCMC sweeps given in Algorithm 1. A single call of Algorithm 1 will transform a current sample $(\theta, Z_{1:n}, \mathbf{X}_{1:n})$ from the posterior to a new sample $(\theta', Z'_{1:n}, \mathbf{X}'_{1:n})$. The entire sequence of samples yielded by the repeated calls to Algorithm 1 will constitute the desired set of Monte Carlo samples from (12). We need though to discard an initial sequence of this set so that the remaining samples retained are correctly distributed. Section 3 is dedicated to the exposition of the first loop of Algorithm 2. The remaining two loops are more easily described. The principal aim of the second loop is to resample the continuous state trajectory of each target using the particle Gibbs sampler [11] (but using the implementation in [22]) which we find enhances our MCMC algorithm's efficiency significantly. This is done by first explicitly isolating the state trajectories of individual targets as in Section 2.4 and then updating the targets' trajectories in turn using the Particle Gibbs sampler. When conjugate priors are available for the components of θ, it is possible to sample $p(\theta | z_{1:n}, \mathbf{x}_{1:n}, y_{1:n}) \propto p(\theta) p_\theta(z_{1:n}, \mathbf{x}_{1:n}, y_{1:n})$ exactly in the final loop, in which case set $n_3 = 1$. Otherwise, one can run a Metropolis-Hastings (MH) algorithm to sample from this pdf. When the MTT parameters are known, the third loop can be omitted and we refer to the resulting algorithm as the *MCMC tracker*.

Algorithm 1: MCMC for state and parameter learning

Input: Current sample $(\theta, z_{1:n}, \mathbf{x}_{1:n})$, data $y_{1:n}$, number of inner loops n_1, n_2, n_3.
Output: Updated sample $(\theta', z'_{1:n}, \mathbf{x}'_{1:n})$.

```
for j = 1 : n_1 do
    Update (z_{1:n}, x_{1:n}) by invoking Algorithm 2.
Isolate target trajectories (see (11)) {(t_b^k, x^k)}_{k=1}^K.
for j = 1 : n_2 do
    for k = 1 : K do
        Update x^k using particle Gibbs conditioned on other trajectories (≠ k) and θ.
Call the updated sample (z'_{1:n}, x'_{1:n}).
for j = 1 : n_3 do
    Conditioned on (z'_{1:n}, x'_{1:n}), update θ to θ' using a MH move or a Gibbs move.
```

3 MCMC moves

In this section, we present the MCMC moves to explore $(Z_{1:n}, \mathbf{X}_{1:n})$ for the first loop in Algorithm 1. Notice that the dimension of $\mathbf{X}_{1:n}$, which is proportional to $\sum_{t=1}^{n} K_t^x$, changes with $Z_{1:n}$. Therefore, the posterior distribution $p_\theta(z_{1:n}, \mathbf{x}_{1:n}|y_{1:n})$ is said to be *trans-dimensional*.

3.1 A brief on trans-dimensional MCMC

A general method for sampling from a trans-dimensional distribution is the reversible jump MCMC (RJMCMC) algorithm of [23]. We briefly describe RJMCMC for the (general) target distribution $\pi(m, x_m)$ where m is a discrete variable (e.g. $m \in \{1, 2 \ldots\}$) known as the model index and $x_m \in \mathbb{R}^{d_m}$. Note though that in general $m' \neq m$ does not imply $d_{m'} \neq d_m$. We now define the procedure for generating samples from a different dimension and either accepting or rejecting them.

For each (m, x_m), let $Q(m'|m, x_m)$ be a probability mass function satisfying $\sum_{m'} Q(m'|m, x) = 1$ and $Q(m'|m, x_m) = 0$ if $d_{m'} = d_m$. Furthermore, for each m' such that $d_{m'} > d_m$, let $Q(u|m, x_m, m')$ be a pdf on $\mathbb{R}^{d_{m'}-d_m}$. Q will be the proposal distribution and $Q(m'|m, x_m) = 0$ if $d_{m'} = d_m$ implies Q only proposes moves across dimension. (Sampling from π can be achieved by combining this proposal with others that do not move across dimensions and then cycling between them either randomly or deterministically. The derivation of the accept-reject probability is routine for intra-dimensional moves.)

Let $(m, x_m) \sim \pi$. First sample m' from $Q(\cdot|m, x_m)$ and if $d_{m'} > d_m$, then sample u from $Q(\cdot|m, x_m, m')$, which are the extra (so called dimension matching) continuous r.v.'s needed to generate the candidate sample $x_{m'} \in \mathbb{R}^{d_{m'}}$. (Indeed it is equivalent to jointly sample (m', u), as opposed to doing it one after the other. The order is application dependent.) Assume $d_{m'} > d_m$. (The reverse case is considered below.) The candidate sample is obtained by applying a *bijection* (to be chosen by the practitioner just as Q was) to yield $x_{m'} = \beta_{m,m'}(x_m, u)$. The acceptance probability for the proposed sample $(m', x_{m'})$ is $\alpha(m', x_{m'}; m, x_m) = \min\{1, r(m', x_{m'}; m, x_m)\}$ where $r(m', x_{m'}; m, x_m)$ is

$$\frac{\pi(m', x_{m'})}{\pi(m, x_m)} \frac{Q(m|m', x_{m'})}{Q(m', u|m, x_m)} |\nabla \beta_{m,m'}(x_m, u)| \tag{13}$$

where the right most term is the Jacobian of $\beta_{m,m'}$. (If $x_{m'} = \beta_{m,m'}(x_m, u)$ is a mapping that permutes the components of the vector (x_m, u) then the Jacobian is 1.) If however $d_{m'} < d_m$,

let $(x'_m, u) = \beta^{-1}_{m',m}(x_m)$ and the candidate sample for the move to the lower dimension model is x'_m. The proposal $(m', x_{m'})$ is accepted with probability $\min\{1, r\}$ where $r(m', x_{m'}; m, x_m)$ is

$$\frac{\pi(m', x_{m'})}{\pi(m, x_m)} \frac{Q(m, u|m', x_{m'})}{Q(m'|m, x_m)} |\nabla \beta_{m',m}(x_{m'}, u)|^{-1} \tag{14}$$

In the MTT model, each target configuration $z_{1:n}$ corresponds to a model index m, $\mathbf{x}_{1:n}$ corresponds to the continuous variable x_m, and $p_\theta(z_{1:n}, \mathbf{x}_{1:n}|y_{1:n})$ corresponds to $\pi(m, x_m)$. The bijections $\beta_{m,m'}$ will do nothing more than permute the input variables to preserve the MTT ordering rule $\mathcal{O}$ and thus all Jacobians are 1.

3.2 MCMC to explore $(Z_{1:n}, \mathbf{X}_{1:n})$ in loop 1 of Algorithm 1

Algorithm 2 proposes a change to $(Z_{1:n}, \mathbf{X}_{1:n})$ by selecting one of the following proposals at random:

1. *birth/death* proposal to create or delete a target;
2. multi-step *extension/reduction* proposal to extend/reduce an existing track by multiple time units;
3. one-step *extension/reduction* proposal to extend/reduce an existing track by one time unit;
4. *state* proposal to exchange states between targets.

All but the state proposal changes the dimension of $\mathbf{X}_{1:n}$ and hence the corresponding RJMCMC acceptance probability in (13) needs to be derived. However the bijections are such that the Jacobian in (13) is always 1. Each proposal type is now described in detail in the subsections below

Algorithm 2: MCMC moves to explore $(Z_{1:n}, \mathbf{X}_{1:n})$

Input: Sample $(z_{1:n}, \mathbf{x}_{1:n})$, data $y_{1:n}$, parameter θ
Output: Updated sample $(z'_{1:n}, \mathbf{x}'_{1:n})$
1 Randomly select a proposal type from
{birth/death, multi-step extension/reduction, state, one-step extension/reduction}.
2 Propose $(z'_{1:n}, \mathbf{x}'_{1:n})$ by executing the chosen proposal.
3 Calculate acceptance prob. $\alpha(z'_{1:n}, \mathbf{x}'_{1:n}; z_{1:n}, \mathbf{x}_{1:n})$ (see (15)/(16), (23)/(24), (26)).
Output $(z'_{1:n}, \mathbf{x}'_{1:n})$ with prob. α, otherwise $(z_{1:n}, \mathbf{x}_{1:n})$.

3.3 Birth/Death Proposal

The birth/death proposal of Algorithm 2 initiates a new track or deletes an existing track and thus only generates proposed samples $(z'_{1:n}, \mathbf{x}'_{1:n})$ (in step 2 which are then tested for acceptance) that move across dimension, i.e. is never intra-dimensional. Birth creates a track sequentially in time, until a stopping rule is met, by using the observed images to increase the probability of acceptance (we describe this sequential process in detail in Sec. 3.4 below,) while death deletes a by choosing one at random.

The current MCMC input sample (Algorithm 2) is $(z_{1:n}, \mathbf{x}_{1:n})$ or $\{(k, t^k_b, \mathbf{x}^k)\}_{k=1}^K$ in the alternative parameterization of 2.4. Sampling from the birth/death proposal commences by

first choosing to create or delete a track with probability 0.5. If birth is chosen, a new track with birth time t_b and states

$$\mathbf{x} = [x_0, \ldots, x_{l-1}], \quad x_i = (a_i, s_i, v_i)$$

are proposed, where (a_i, s_i, v_i) denote the intensity, spatial coordinates and velocity components of the state at time $t = t_b + i$. If death is chosen, one of the K targets are randomly deleted, say target k with probability $q_{d,\theta}(k|z_{1:n}, \mathbf{x}_{1:n}, y_{1:n})$. In the case of a birth, using $(z_{1:n}, \mathbf{x}_{1:n})$ and the newly created track, the ordering rule of Section 2.4 is invoked to obtain the MTT proposed sample $(z'_{1:n}, \mathbf{x}'_{1:n})$. Assume the newly created target has label k' in the alternative parameterization of $(z'_{1:n}, \mathbf{x}'_{1:n})$. The acceptance probability is $\alpha_1 = \min\{1, r_1\}$ where $r_1(z'_{1:n}, \mathbf{x}'_{1:n}; z_{1:n}, \mathbf{x}_{1:n})$ is

$$\frac{p_\theta(z'_{1:n}, \mathbf{x}'_{1:n}, y_{1:n})}{p_\theta(z_{1:n}, \mathbf{x}_{1:n}, y_{1:n})} \frac{q_{d,\theta}(k'|z'_{1:n}, \mathbf{x}'_{1:n}, y_{1:n})}{q_{b,\theta}(t_b, \mathbf{x}|z_{1:n}, \mathbf{x}_{1:n}, y_{1:n})} \tag{15}$$

The term $q_{b,\theta}(t_b, \mathbf{x}|z_{1:n}, \mathbf{x}_{1:n}, y_{1:n})$, defined in Sec. 3.4 below, is the pdf of the newly created target states (which corresponds to term $Q(m', u|m, x_m)$ in the denominator of (13).) The term $q_{d,\theta}(k'|z'_{1:n}, \mathbf{x}'_{1:n}, y_{1:n})$ is the probability of deleting this newly created target.

If death is chosen, delete target k of $\{(i, t_b^i, \mathbf{x}^i)\}_{i=1}^K$ with probability $q_{d,\theta}(k|z_{1:n}, \mathbf{x}_{1:n}, y_{1:n})$ and let $(z'_{1:n}, \mathbf{x}'_{1:n})$ be the new MTT state excluding target k. The acceptance probability is $\alpha_1 = \min\{1, r_1\}$ where $r_1(z'_{1:n}, \mathbf{x}'_{1:n}; z_{1:n}, \mathbf{x}_{1:n})$ is

$$\frac{p_\theta(z'_{1:n}, \mathbf{x}'_{1:n}, y_{1:n})}{p_\theta(z_{1:n}, \mathbf{x}_{1:n}, y_{1:n})} \frac{q_{b,\theta}(t_b^k, \mathbf{x}^k|z'_{1:n}, \mathbf{x}'_{1:n}, y_{1:n})}{q_{d,\theta}(k|z_{1:n}, \mathbf{x}_{1:n}, y_{1:n})} \tag{16}$$

3.4 Birth/Death Proposal: Creating a new target

The birth procedure below proposes a new track by using the *residual* images to first construct the intensity and spatial coordinates of the track and then finally the velocity values of the track. For the current MCMC input sample (Algorithm 2) $(z_{1:n}, \mathbf{x}_{1:n})$, subtract the contribution of the k_t^x targets and background intensity b_t from the image y_t to get the residual image y_t^r where

$$y_{t,i}^r = y_{t,i} - h_i(\mathbf{x}_t) - b_t, \qquad i = 1, \ldots, m. \tag{17}$$

Match filter y_t^r to get the image y_t^f where the j-th pixel in the filtered image is

$$y_{t,j}^f = \frac{1}{E_{\bar{h}}} \sum_{i=1}^m y_{t,i}^r \bar{h}_i(j) \tag{18}$$

where $\bar{h}_i(j)$ is defined in (3)-(4) and $E_{\bar{h}} = \sum_{i=1}^m \bar{h}_i(j)^2$ is the energy of the filter $\{\bar{h}_i(j)\}_{i=1}^m$. (The sum that defines $y_{t,j}^f$ can be truncated to $i \in L(j)$.) The rationale is that the presence of a target at or close to pixel j will likely result in $y_{t,j}^f$ being a local maxima among pixels. Put another way, local maxima of $y_{t,j}^f$ are likely locations of targets.

3.4.1 Proposing the initial state

The move is commenced by choosing the birth time $t = t_b$ randomly from $1, \ldots, n$ and then followed by steps 1 and 2 below.

Step 1 Let

$$G_t = \{1 \leq i \leq m : y^f_{t,i} \text{ is a local maxima}, y^f_{t,i} \geq \gamma_t(\theta)\}$$

and randomly choose $i \in G_t$. $\gamma_t(\theta)$ is a time-dependent threshold chosen to avoid peaks that are not likely to be target generated. A definition is given in the numerical section (Sec. 4.)

As a local intensity maxima is not necessarily target generated, perform a hypothesis test on the square of $l \times l$ pixels $L(i)$ centered at chosen maxima $i \in G_t$. Let $y^r_{t,L(i)} = \{y^r_{t,j}, j \in L(i)\}$, H_1 the hypothesis that $y^r_{t,L(i)}$ is generated by a new born target and H_0 the converse that it is purely background noise generated. Calculate the test ratio

$$\rho(y^r_{t,L(i)}) = \frac{p(H_1)}{p(H_0)} \frac{p(y^r_{t,L(i)}|H_1)}{p(y^r_{t,L(i)}|H_0)} \tag{19}$$

While $p(y^r_{t,L(i)}|H_0) := \prod_{j \in L(i)} \mathcal{N}(y^r_{t,j}; 0, \sigma^2_{r,t})$ can be calculated analytically, $p(y^r_{t,L(i)}|H_1)$ is intractable but can be estimated, e.g. we use the Laplace approximation. (See Appendix A for details.) The probability of H_1 is $p(H_1) = \sum_{k>k^t_b} \mathcal{P}(k, \lambda_b)$, which is the probability that the number of births exceeds k^b_t where k^b_t is the number of targets born at time t in the current MCMC sample $z_{1:n}$. Set $p(H_0) = 1 - p(H_1)$. H_1 is accepted with probability $\min\{1, \rho(y^r_{t,L(i)})\}$. Return (exit birth move) if H_1 is rejected. (Alternatively, H_1 could be accepted with probability $\rho/(1+\rho)$ which is less presumptuous than $\min\{1, \rho\}$.)

Step 2 After accepting H_1, sample the intensity and position components of the initial target state by

$$(A_0, S_0) \sim \mathcal{N}_{t,i}(\cdot),$$

where $\mathcal{N}_{t,i}(\cdot)$ is a Gaussian derived from the Laplace approximation of the hypothesis test and is given in Appendix A. The subscript (t, i) indicates this approximation is specific to pixels $L(i)$ of time t.

3.4.2 Proposing the remaining intensity-position trajectory

The birth move continues for $t = t_b + k$, $k > 0$, using steps 1 and 2 above until a stopping rule is met. To conserve computations and increase its effectiveness, the range $1 \leq i \leq m$ in the definition of G_t is limited to a region of pixels R_t where the next state would almost surely lie in. For example, R_t can be determined by the previous position s_{t-1} and the upper limit on the target velocity.

The birth move stops at some time $t = t_d$, yielding a target lifespan $l = t_d - t_b$, when either $t > n$, G_t is empty or H_1 is rejected. The output of this move is $(a_1, s_1), \ldots, (a_{l-1}, s_{l-1})$ where (a_0, s_0) was generated before.

3.4.3 Proposing the velocity trajectory

The output of the birth move thus far is $(a_0, s_0), \ldots, (a_{l-1}, s_{l-1})$ which is the complete trajectory of intensity and positions of the new born target. The velocity components are now generated to yield

$$\mathbf{x} = [x_0, \ldots, x_{l-1}], \quad x_i = (a_i, s_i, v_i).$$

For linear Gaussian state dynamics, the velocity can be sampled conditioned on the spatial locations $s_0, \ldots, s_{l-1}$ and more generally, a Gaussian approximation/quadrature technique could be employed.[2]

[2]The numerical examples (synthetic and real data) assume Gaussian targets.

3.4.4 Proposal density of the birth move

Denoting $t_k = t_b + k$, $k = 0, \ldots, l$, we can write the pdf for proposing $(t_b, \mathbf{x})$ in the birth move as

$$\begin{aligned} q_{b,\theta}(t_b, \mathbf{x}|z_{1:n}, \mathbf{x}_{1:n}, y_{1:n}) &= q_0(t_b) q_1(a_0, s_0|y^r_{t_0}) \\ &\times \prod_{k=1}^{l-1} q_2(a_k, s_k|a_{0:k-1}, s_{0:k-1}, y^r_{t_k}) \\ &\times q_3(\text{stop}|a_{0:l-1}, s_{0:l-1}, y^r_{t_l})\ q_4(v_{0:l-1}|a_{0:l-1}, s_{0:l-1}) \end{aligned} \tag{20}$$

where $q_0(t_b)$ is the probability of choosing the birth time t_b, q_1 corresponds to proposing the target's initial intensity and position, and can be written as

$$\begin{aligned} q_1(a_0, s_0|y^r_{t_0}) &= \mathbb{I}[G_{t_0} \neq \emptyset] \\ &\times \sum_{i \in G_{t_0}} q(i|G_{t_0}) \min(1, \rho(y^r_{t_0, L(i)})) \mathcal{N}_{t_0,i}(a_0, s_0) \end{aligned} \tag{21}$$

where $\mathbb{I}[G_{t_0} \neq \emptyset]$ is 1 if G_{t_0} is non-empty; $q(i|G_{t_0})$ is the probability of choosing local intensity peak $i \in G_{t_0}$; $\min(1, \rho(y^r_{t_0,L(i)}))$ is the probability of accepting hypothesis H_1 at peak $i \in G_{t_0}$; $\mathcal{N}_{t_0,i}$ is the Gaussian density approximation for the initial value.

The law q_2 is that for adding more states sequentially and may be written as

$$\begin{aligned} q_2(a_k, s_k|a_{0:k-1}, s_{0:k-1}, y^r_{t_k}) &= \mathbb{I}[G_{t_k} \neq \emptyset] \\ &\times \sum_{i \in G_{t_k}} q(i|G_{t_k}) \min(1, \rho(y^r_{t_k, L(i)})) \mathcal{N}_{t_k,i}(a_k, s_k) \end{aligned} \tag{22}$$

which is similar to q_1 except that the previously created states $a_{0:k-1}$ and $s_{0:k-1}$ are used to calculate $\rho(y^r_{t_k,L(i)})$, which is defined as in (19) with the difference here being $p(H_1) = p_s$ (target survival probability.) The likelihood ratio term in (19) is derived in the Appendix, as is the Gaussian term $\mathcal{N}_{t_k,i}$ in (22).

When $t_l = t_b + l > n$ stopping at t_l is certain. Otherwise, the law q_3 corresponds to stopping due to the target not surviving, G_{t_l} being empty, or the hypothesis test failing and is given by

$$\begin{aligned} &q_3(\text{stop}|a_{0:l-1}, s_{0:l-1}, y^r_{t_l}) \\ &= 1 - p_s \mathbb{I}[G_{t_l} \neq \emptyset] \sum_{i \in G_{t_l}} q(i|G_{t_l}) \min(1, \rho(y^r_{t_l, L(i)})) \end{aligned}$$

For linear and Gaussian state dynamics, conditioned on $a_0, s_0, \ldots, a_{l-1}, s_{l-1}$, the velocity can be sampled (exactly) since q_4 will be a Gaussian distribution as well.

3.5 Multi-step Extension/Reduction proposal

This proposal extends or reduces the trajectory of a randomly chosen target. A target's trajectory is extended (/reduced) by bringing forward its birth (/death) time or delaying its death (/birth) time. The extra state values are then appended (/discarded) accordingly. Like the birth/death proposal, this proposal only moves the MCMC sample across dimension.

The current MCMC input sample (Algorithm 2) is $(z_{1:n}, \mathbf{x}_{1:n})$ or $\{(k, t^k_b, \mathbf{x}^k)\}_{k=1}^K$ in the alternative parameterization of 2.4. Sampling from the multi-step extension/reduction proposal commences with choosing between extension and reduction equiprobably.

Extension From the subset of targets *with* lifetimes less than n, randomly select a target and extension direction. (Without the lifetime restriction the chosen target cannot be extended further.) The direction of extension is chosen equiprobably if both forward and backward extensions are permissible. Assume target k is selected for a forward extension, denoted k_+. A new (delayed) death time and trajectory extension $\mathbf{x} = (x_1, x_2, \ldots)$ is proposed to yield the new extended trajectory $\hat{\mathbf{x}}^k = (\mathbf{x}^k, \mathbf{x})$. For a backward extension of target k, the event denoted by k_-, a new (earlier) birth time, denoted τ_b^k and trajectory extension $\mathbf{x}$ is proposed to yield $\hat{\mathbf{x}}^k = (\mathbf{x}, \mathbf{x}^k)$. The ordering rule of Section 2.4 is invoked to obtain the MTT proposed sample $(z'_{1:n}, \mathbf{x}'_{1:n})$ from the unaltered targets $\{(t_b^i, \mathbf{x}^i)\}_{i=1, i\neq k}^K$ and the extended target $(\tau_b^k, \hat{\mathbf{x}}^k)$. The acceptance probability is $\alpha_2 = \min\{1, r_2\}$ where $r_2(z'_{1:n}, \mathbf{x}'_{1:n}; z_{1:n}, \mathbf{x}_{1:n})$ is

$$\frac{p_\theta(z'_{1:n}, \mathbf{x}'_{1:n}, y_{1:n})}{p_\theta(z_{1:n}, \mathbf{x}_{1:n}, y_{1:n})} \frac{q_{r,\theta}(z_{1:n}, \mathbf{x}_{1:n} | z'_{1:n}, \mathbf{x}'_{1:n})}{q_{e,\theta}(k_{+/-}, \mathbf{x} | z_{1:n}, \mathbf{x}_{1:n}, y_{1:n})}. \tag{23}$$

The probability density of choosing target k, extension direction and states $\mathbf{x}$ is denoted $q_{e,\theta}(k_{+/-}, \mathbf{x} | z_{1:n}, \mathbf{x}_{1:n}, y_{1:n})$ which is calculated similarly to the expression (20) of the birth move and it is not repeated here. We denote the total probability of making the return transition from $(z'_{1:n}, \mathbf{x}'_{1:n})$ to $(z_{1:n}, \mathbf{x}_{1:n})$ via the reduction, described next, with the term $q_{r,\theta}$ above.

Reduction Randomly select a target from the subset of targets with lifetimes exceeding one and then the reduction direction, either forwards or backwards, equiprobably. Let k_+ denote target k for forward reduction and k_- for backward reduction. In the forward case, a reduction time point $t \in \{t_b^k + 1, \ldots, t_d^k - 1\}$ is chosen randomly and discard the time $\{t, \ldots, t_d^k - 1\}$ section of the track, causing t to be the new death time. If backwards then $t \in \{t_b^k, \ldots, t_d^k - 2\}$ is chosen randomly and the $\{t_b^k, \ldots, t\}$ portion is discarded to yield a new birth time of $t+1$. Let τ_b^k denote the (possibly new) birth time, $\hat{\mathbf{x}}^k$ the retained trajectory and $\mathbf{x}$ the discarded forward/backward state trajectory of target k. The ordering rule of Section 2.4 is invoked to obtain the MTT proposed sample $(z'_{1:n}, \mathbf{x}'_{1:n})$ from the the unaltered targets $\{(t_b^i, \mathbf{x}^i)\}_{i=1, i\neq k}^K$ and the reduced target $(\tau_b^k, \hat{\mathbf{x}}^k)$. Assume the reduced target has label k' in the alternative parameterization of $(z'_{1:n}, \mathbf{x}'_{1:n})$. Let $q_{r,\theta}(z'_{1:n}, \mathbf{x}'_{1:n} | z_{1:n}, \mathbf{x}_{1:n})$ denote the total probability of making the transition from $(z_{1:n}, \mathbf{x}_{1:n})$ to $(z'_{1:n}, \mathbf{x}'_{1:n})$ via the described reduction step. The acceptance probability is $\alpha_2 = \min\{1, r_2\}$ where $r_2(z'_{1:n}, \mathbf{x}'_{1:n}; z_{1:n}, \mathbf{x}_{1:n})$ is

$$\frac{p_\theta(z'_{1:n}, \mathbf{x}'_{1:n}, y_{1:n})}{p_\theta(z_{1:n}, \mathbf{x}_{1:n}, y_{1:n})} \frac{q_{e,\theta}(k'_{+/-}, \mathbf{x} | z'_{1:n}, \mathbf{x}'_{1:n}, y_{1:n})}{q_{r,\theta}(z'_{1:n}, \mathbf{x}'_{1:n} | z_{1:n}, \mathbf{x}_{1:n})} \tag{24}$$

3.6 One-step Extension/Reduction proposal

The intensity threshold $\gamma_t(\theta)$ used in the Birth (Sec. 3.3)) and Extension (Sec. 3.5) moves ignore the local intensity peaks G_t of the match filtered image y_t^f that are below $\gamma_t(\theta)$. This may result in the next state of a momentarily dim target not being detected. As a remedy, a new *one-step* extension/reduction proposal is defined. This proposal, which is to be regarded as a different proposal to multi-step extension/reduction (Sec. 3.5), proceeds as multi-step except that it extends or truncates the trajectory of the selected target by one time point only. In particular, its acceptance probability is $\alpha_3 = \min\{1, r_3\}$ with r_3 defined as in (23) and (24) but with the following differences. The expression $q_{r,\theta}(z'_{1:n}, \mathbf{x}'_{1:n} | z_{1:n}, \mathbf{x}_{1:n})$ in (23) (and (24)) is now $1/2K$ (assuming K targets have lifetimes exceeding one) since the probability of selecting a particular target is $1/K$ and then the reduction direction is $1/2$. Let the trajectory of the selected target prior to extension be $\mathbf{x}^k = (x_0^k, \ldots, x_{l^k-1}^k)$, then $q_{e,\theta}$ in (23) for k_+ is

$$q_{e,\theta}(k_+, x | z_{1:n}, \mathbf{x}_{1:n}) = q_{e,\theta}(k_+ | z_{1:n}, \mathbf{x}_{1:n}) f_\psi(x | x_{l^k-1}^k)$$

where the first factor is the probability of selecting k_+ (target k and forward extension) and the extended state value is sampled from the prior model f_ψ. For a backward extension or k_-, $q_{e,\theta}(k_-, x|z_{1:n}, \mathbf{x}_{1:n})$ is

$$q_{e,\theta}(k_-|z_{1:n}, \mathbf{x}_{1:n}) \frac{\mu_\psi(x) f_\psi(x_0^k|x)}{\int \mu_\psi(x') f_\psi(x_0^k|x') \mathrm{d}x'},$$

the extended state is sampled from initial distribution μ_ψ of new targets conditioned on the value of its next state.

3.7 State proposal

This proposal chooses a pair of targets and swaps a section of their state trajectories. In particular, it randomly chooses a time $t < n$ and then randomly changes I_{t+1}, which is the vector that links targets in $\mathbf{X}_t$ and $\mathbf{X}_{t+1}$, as illustrated in Figure 2. When $X_{t,i}$ has descendant $X_{t+1,g}$, it can propose to swap its descendant with that of $X_{t,j}$ (case 1), or change its descendant to the initial state $X_{t+1,h}$ of a target born at time $t+1$ (case 2), or to delete the link (case 4). When $X_{t,i}$ has no descendant, it can be merged with a new-born target at time $t+1$ by linking to its initial state (case 3), or steal another surviving target's descendant (case 5). The state proposal is purely intra-dimensional (or in the context of Sec. 3.1 it moves between models pairs (m, m') satisfying $d_m = d_{m'}$.)

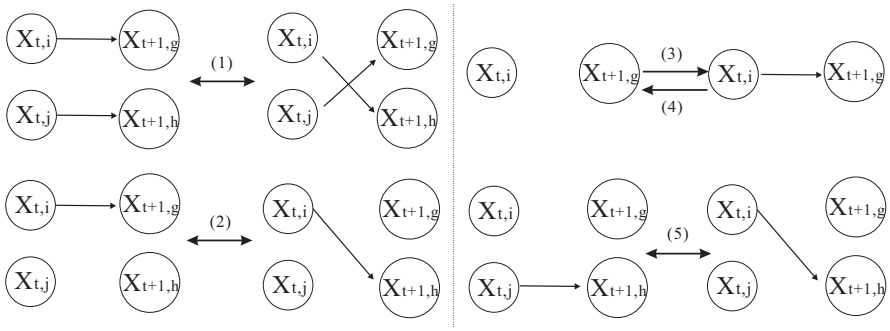

Figure 2: State move.

The current MCMC input sample (Algorithm 2) is $(z_{1:n}, \mathbf{x}_{1:n})$ or $\{(k, t_b^k, \mathbf{x}^k)\}_{k=1}^K$ in the alternative parameterization of 2.4. Sampling from the state proposal commences with choosing a time $t < n$ and a pair of targets (labels) $U = \{i, j\}$ from the total set of targets $\{1, \ldots, K\}$ subject to targets i and j being alive at times t and $t+1$ respectively. Choosing $i = j$ is permitted (then $U = \{i\}$) and target i must be alive at time t and $t+1$. (For example, in implementation we chose a target i from time t, and a target state value from the set of time $t+1$ targets with probability inversely proportional to the distance from target i's state value at time t.) We denote the probability of a particular selection by $q_{s,\theta}(t, U|\mathbf{x}_{1:n}, z_{1:n})$. If $i = j$ then split target i into two targets as in case 4 of Figure 2. If $i \neq j$, pair the *ancestors* of i with the *descendants* of j in the manner shown in Figure 2 (all cases except 4), i.e. the swap alters trajectories $\mathbf{x}^i$ and $\mathbf{x}^j$ to

$$\begin{aligned} \mathbf{x}^i &\rightarrow (x_0^i, \ldots, x_s^i, x_{s'}^j, \ldots, x_{l^j-1}^j) =: \hat{\mathbf{x}}^i \\ \mathbf{x}^j &\rightarrow (x_0^j, \ldots, x_s^j, x_{s'}^i, \ldots, x_{l^i-1}^i) =: \hat{\mathbf{x}}^j \end{aligned} \tag{25}$$

where $t_b^i + s = t$, $t_b^j + s' = t+1$. Note the birth time (of at most one target) may change. Call the birth times after the swap τ_b^i, τ_b^j

The next part of the state move then proposes a change to the continuous variables of the affected targets $k \in U = \{i, j\}$ form $\hat{\mathbf{x}}^k$ to $\tilde{\mathbf{x}}^k$ to increase the chance of the move being accepted. This is because changing the links between targets at time t and $t+1$ can cause a mismatch in the velocity and intensity of the newly formed links. As such, the state move will then propose a change to the velocity and intensity components of the affected targets while retaining their original spatial position components. Let $q_{s,\theta}(t, U, (\tilde{\mathbf{x}}^k)_{k\in U}|\mathbf{x}_{1:n}, z_{1:n}, y_{1:n})$ denote the joint pdf of selecting (t, U) and the change $(\hat{\mathbf{x}}^i, \hat{\mathbf{x}}^j) \to (\tilde{\mathbf{x}}^i, \tilde{\mathbf{x}}^j)$. (See Appendix B for the expression.) Finally, we let $(z'_{1:n}, \mathbf{x}'_{1:n})$ denote MTT state (in the original parameterization) of the unaltered $\{t_b^k, \mathbf{x}^k\}_{k\in\{1\ldots,K\}/U}$ and altered targets $\{\tau_b^k, \tilde{\mathbf{x}}^k\}_{k\in U}$. The acceptance probability of the state move is then $\alpha_3 = \min\{1, r_3\}$ where $r_3(z'_{1:n}, \mathbf{x}'_{1:n}; z_{1:n}, \mathbf{x}_{1:n})$ is

$$\frac{p_\theta(z'_{1:n}, \mathbf{x}'_{1:n}, y_{1:n})}{p_\theta(z_{1:n}, \mathbf{x}_{1:n}, y_{1:n})} \frac{q_{s,\theta}(t, U', (\mathbf{x}^k)_{k\in U'}|\mathbf{x}'_{1:n}, z'_{1:n}, y_{1:n})}{q_{s,\theta}(t, U, (\tilde{\mathbf{x}}^k)_{k\in U}|\mathbf{x}_{1:n}, z_{1:n}, y_{1:n})} \tag{26}$$

U' is the label set of the targets with swapped components in the alternate parameterisation of $(z'_{1:n}, \mathbf{x}'_{1:n})$.

4 Numerical examples

This section presents the two main numerical examples. The first one uses synthetic data and assumes known MTT parameters so that a fair comparison can be made between our MCMC tracker (Algortihm 1 excluding parameter learning) and the multi-Bernoulli (MB) filter of [9]. (The MB tracker does not learn parameters.) The numerical results will demonstrate the performance improvements of our method when tracking targets that are close to each other with overlapping illumination regions. The second example is a real-data example that applies Algorithm 1 to track Fab[3] labelled Jurkat T-cells. The tracking method currently used by biochemists [2] extracts point measurements from the images and then connects them to form trajectories using a nearest neighbour type method. Our algorithm will be shown to outperform theirs when tracking dim targets as well as targets with overlapping illumination regions. All simulations were run in Matlab on a PC with an Intel i5 2.8 GHZ $\times 2$ processor.

Recall the definition of an individual target's state $X_t = (A_t, S_t(1), S_t(2), V_t(1), V_t(2))$ in (1). For the numerical examples, we use a drifting intensity and near constant velocity motion model,

$$\begin{aligned} A_t = A_{t-1} + U_t, \quad U_t &\overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_i^2), \\ [S_t(j), V_t(j)] &= [S_{t-1}(j) + \delta V_{t-1}(j), V_{t-1}(j)] + U_{t,j}, \\ U_{t,1}^{\mathrm{T}} &\overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_x^2\Sigma), U_{t,2}^{\mathrm{T}} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_y^2\Sigma), \end{aligned}$$

where δ is the (known) sampling interval and

$$\Sigma = \begin{pmatrix} \delta^3/3 & \delta^2/2 \\ \delta^2/2 & \delta \end{pmatrix}.$$

The initial hidden state is assumed to be Gaussian distributed with mean $\mu_b = (\mu_{bi}, \mu_{bx}, \mu_{by}, 0, 0)^{\mathrm{T}}$ and covariance $\Sigma_b = \mathrm{diag}(\sigma_{bi}^2, \sigma_{bp}^2, \sigma_{bp}^2, \sigma_{bv}^2, \sigma_{bv}^2)$. The mean of the initial velocity is set to be 0 in the absence of more information but this still can yield directional motion if the observations support this. All the parameters ψ of the hidden state dynamics are (see (1))

$$\psi = (\mu_{bi}, \mu_{bx}, \mu_{by}, \sigma_{bi}^2, \sigma_{bp}^2, \sigma_{bv}^2, \sigma_i^2, \sigma_x^2, \sigma_y^2).$$

[3] Fab (Fragment antigen-binding) is a region on an antibody that binds to antigens.

and augmenting ψ with the parameters of the target birth/death and observation models gives

$$\theta = (\psi, p_s, \lambda_b, b_1, \sigma^2_{r,1}, \ldots, b_n, \sigma^2_{r,n}).$$

Prior for θ: All the variance components above have independent priors, which is the inverse gamma distribution $\mathcal{IG}(\alpha_0, \beta_0)$ with (common) shape α_0 and scale β_0 parameters. (Setting $\alpha_0 \ll 1$ and $\beta_0 \ll 1$ yields a less informative prior.) Given σ^2_{bi}, σ^2_{bpx}, σ^2_{bpy} and $\sigma^2_{r,t}$ (for $t = 1, \ldots, n$), the priors of μ_{bi}, μ_{bx}, μ_{by} and b_t are $\mu_{bi}|\sigma^2_{bi} \sim \mathcal{N}(\mu_0, \sigma^2_{bi}/n_0)$, $\mu_{bx}|\sigma^2_{bp} \sim \mathcal{N}(\mu_0, \sigma^2_{bp}/n_0)$, $\mu_{by}|\sigma^2_{bp} \sim \mathcal{N}(\mu_0, \sigma^2_{bp}/n_0)$, $b_t|\sigma^2_{r,t} \sim \mathcal{N}(\mu_0, \sigma^2_{r,t}/n_0)$. These are made more diffused by setting n_0 and μ_0 to be small. The conjugate priors of p_s, λ_b are

$$p_s \sim \text{Unif}(0,1), \quad \lambda_b \sim \mathcal{G}(\alpha_0, \beta_0),$$

where $\text{Unif}(a,b)$ and $\mathcal{G}(\alpha, \beta)$ represent (respectively) the uniform distribution over (a,b) and the gamma distribution with shape parameter α and scale parameter β. We set $\alpha_0 \ll 1, \beta_0 \gg 1$ to make the prior less informative. (If prior knowledge is available, a more appropriate choice of (α_0, β_0) can be made.) K^s_t is a Binomial r.v. with success probability p_s and number of trials K^x_{t-1}. K^b_t is a Poisson r.v. with rate λ_b. Thus their posteriors are

$$p_s|z_{1:n}, y_{1:n} \sim \mathcal{B}\bigg(1 + \sum_{t=1}^{n} k^s_t,\ 1 + \sum_{t=2}^{n}(k^x_{t-1} - k^s_t)\bigg),$$
$$\lambda_b|z_{1:n}, y_{1:n} \sim \mathcal{G}\bigg(\alpha_0 + \sum_{t=1}^{n} k^b_t,\ (\beta_0^{-1} + n)^{-1}\bigg).$$

The illuminated region $L(s)$ is an $l \times l$ square region of pixels centered at s, with $l = 1 + \lceil 4\sigma_h \rceil / \Delta$ where $\lceil \cdot \rceil$ rounds up its argument. The intensity threshold $\gamma_t(\theta)$ is chosen to be $\gamma_t(\theta) = \min(\mu_{bi} - 3\sigma_{bi}, 3\sigma_{r,t}/\sqrt{E_{\bar{h}}})$ using the following rationale. We expect $y^f_{t,j}$ in (18) to exceed $\mu_{bi} - 3\sigma_{b,i}$ (mean birth illumination minus 3 times its standard deviation) with high probability if a target is present in pixel j of the residual image in (17). However, assuming no targets illuminate pixels $L(j)$ of the residual image, $\sigma_{r,t}/\sqrt{E_{\bar{h}}}$ is the standard deviation of $y^f_{t,j}$. With high probability, $y^f_{t,j}$ should not be exceed $3\sigma_{r,t}/\sqrt{E_{\bar{h}}}$ and avoids triggering detection.

4.1 Comparison with the multi-Bernoulli tracker

We compared our algorithm with the MB tracker [9]. Unlike the subsequent real data example, this synthetic case assumed $b_t = 0$ and $\sigma^2_{r,t} = \sigma^2_r$ for all t (see (6).) We synthesised 50 frames (images) of 168×184 pixels each using the parameter vector

$$\psi^* = (30, 0, 0, 4, 25, 3, 0.5, 0.3, 0.7), \theta^* = (\psi^*, 0.95, 0.3, 1). \tag{27}$$

We set $\sigma^2_h = 1$, $\Delta = 1$. This gives a 5×5 pixel square for the illuminated region $L(s)$ (see (3)). From (3), define $\text{SNR} = 20\log(\frac{a\Delta^2/2\pi\sigma^2_h}{\sigma_r})$. For $a = \mu_{bi} = 30$, the initial SNR is 13.6 dB. The synthetic data had 20 targets whose trajectories are shown in Figure 3a along with the trajectories obtained by running the MCMC tracker with $n_1 = 30$, $n_2 = 1$ (and $n_3 = 0$) and 15 particles per target for the particle Gibbs sampler. Figure 3a shows all targets being tracked completely. In contrast, Figure 3b shows the output of the MB tracker of [9]. The birth process assumed by the MB tracker has four terms each of which has the same initial distribution $\mathcal{N}(\cdot|\mu_b, \Sigma_b)$ and existence probability 0.1. Pruning and merging targets are performed as suggested in [9] to eliminate tracks with existence probability less than 0.01 and

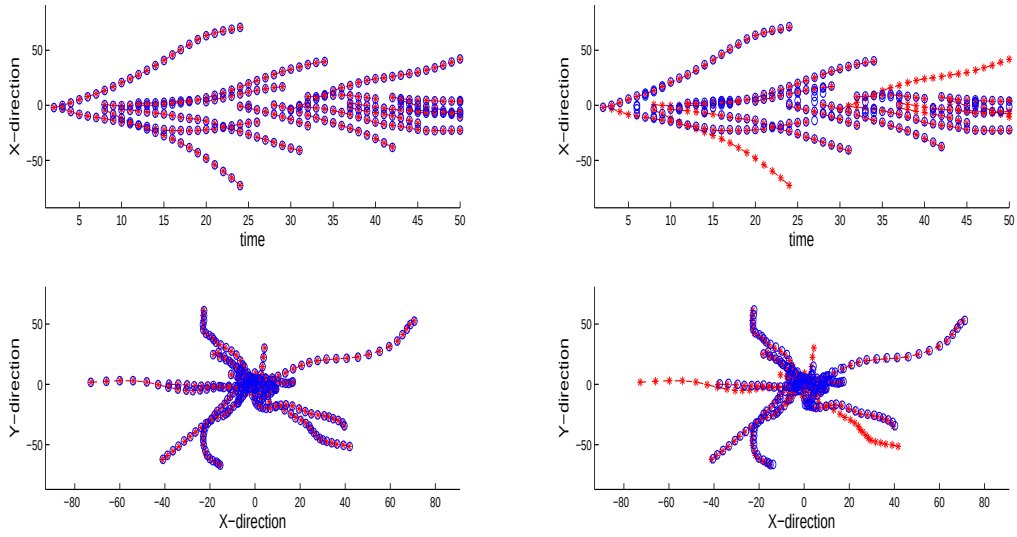

(a) MCMC tracker **(b)** multi-Bernoulli tracker

Figure 3: Comparison with the multi-Bernoulli filter in [9]. Tracks labelled $-*-$ (red) is ground truth while (blue) circles are the estimates.

merge two tracks when they fall within a fraction (3/4) of a pixel size in distance. The number of particles assigned for each hypothesised target in MB tracker is restricted between 5000 and 8000. In Figure 3b, it is seen that some tracks are lost after they cross, which is the main limitation of the MB tracker as pointed out in [9]. This is because crossing targets invalidates the crucial assumption, necessary to derive the MB tracker, that the illuminated region of the targets do not overlap. In terms of the computation time, the MB tracker take less than one minute to run while the MCMC tracker takes 6 minutes. Closely spaced targets are common place in many applications, an example being the real-data experiment reported below. Our MCMC tracker should be viewed as a method applicable to closely spaced targets and not as a competitor of a technique optimized for non-overlapping targets like the MB tracker.

The previous comparison was done assuming known model parameters. The current example revisits the same data set assuming θ^* in (27) is unknown to Algorithm 1, which was re-run with the initial parameter set to

$$\theta^{(0)} = (45, 10, 5, 8, 50, 6, 3, 1, 1.5, 0.6, 1, 4),$$

while MB output of the previous example was used. (The MB was given the true parameter as it does not incorporated parameter learning. The MCMC outputs of Algorithm 1 will be denoted $(z_{1:n}^{(i)}, \mathbf{x}_{1:n}^{(i)}, \theta^{(i)})$.) OSPA distances [9] of the three algorithms are plotted in Figure 4a. Figure 4a shows that the tracking performance with unknown parameters is similar to the known case reported earlier. A further verification is the probability density values plotted in Figure 4b where $p_{\theta*}(z_{1:n}^{(i)}, \mathbf{x}_{1:n}^{(i)}, y_{1:n})$ was calculated from the previous experiment (Algorithm 1 with known parameters) and $p_{\theta^{(i)}}(z_{1:n}^{(i)}, \mathbf{x}_{1:n}^{(i)}, y_{1:n})$ are the density values from Algorithm 1 with parameter learning.

Figure 5 shows the histograms of 2000 post burn-in parameter samples of Algorithm 1 as the approximation of $p(\theta|y_{1:n})$. The (red) dashed lines show the MLE estimate $\hat{\theta}$ obtained using

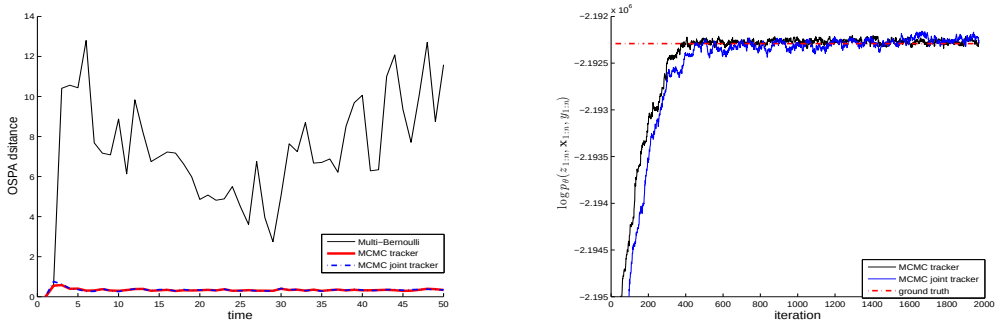

(a) OSPA **(b)** Log density

Figure 4: **(a)** Comparing OSPA tracking error of the multi-Bernoulli filter (top solid black line) with Alg. 1 with and without parameter learning (lower traces.) OSPA error of MCMC almost equal when θ^* is known or learnt during tracking. **(b)** Plot of $p_{\theta*}(z_{1:n}^{(i)}, \mathbf{x}_{1:n}^{(i)}, y_{1:n})$ (Alg. 1 with known θ^*) and $p_{\theta^{(i)}}(z_{1:n}^{(i)}, \mathbf{x}_{1:n}^{(i)}, y_{1:n})$ (Alg. 1 with parameter learning) against MCMC iteration i. Horizontal (red) line indicates ground truth $p_{\theta^*}(z_{1:n}^*, \mathbf{x}_{1:n}^*, y_{1:n})$.

the true value of the latent variables, i.e. $(z_{1:n}^*, \mathbf{x}_{1:n}^*)$. Specifically, $\hat{\theta}$ is comprised of (see (8), (9), (10))

$$(\hat{p}_s, \hat{\lambda}_b) = \arg\max_{p_s, \lambda_b} p(z_{1:n}^*), \quad \hat{\psi} = \arg\max_{\psi} p(\mathbf{x}_{1:n}^* | z_{1:n}^*),$$
$$\hat{\sigma}_r = \arg\max_{\sigma_r} p(y_{1:n} | \mathbf{x}_{1:n}^*, z_{1:n}^*). \tag{28}$$

As a correctness check, for an uninformative prior, the posterior modes should be consistent with MLE of θ^*. The true MLE is $\arg\max_\theta p_\theta(y_{1:n})$, which will be different from (27), is not available as $p_\theta(y_{1:n})$ is intractable. We use $\hat{\theta}$ of (28) as the surrogate. Note the modes of the posterior do coincide with the surrogate MLE.

4.2 Experiments on Fab labelled Jurkat T-cells

The source of the data were Jurkat T-cells, an immortalised cell line of human T-lymphocytes which plays an important function in immune response. Cells were imaged using a microscopy technique called *total internal reflection fluorescence microscopy*. The cell's molecules of interest were bound to antibodies labelled with a bright green-fluorescent dye (Alexafluo48). It is known that the number of labeled molecules (or 'targets') can be high; at physiological levels there can be several molecules per square micron.

The data is comprised of 20 frames of 115×120 pixel images with a pixel size $176nm$ and a frame rate of 17.8 frame/s. The diffusion coefficient D is expected to be in the region of 0.01-$0.1um^2/s$. This implies the displacement of molecules between consecutive frames is expected to be in the order of 0.1-1 pixel. The estimated initial SNR here is around 10dB (estimated from the tracking result). Figure 7 (left) shows one frame of the observed images.

4.2.1 Parameter initialisation

In choosing the initial parameter vector $\theta^{(0)}$, some components were chosen arbitrarily while others were guided by the observed images. (Note that the initialisation step does not need to be overly precise as our algorithm does not depend on specific initial values to work.) We set $p_s = 0.6, \lambda_b = 0.2$ arbitrarily; $\mu_{bx}^{(0)} = \mu_{by}^{(0)} = 60$ to coincide with the image centres since the

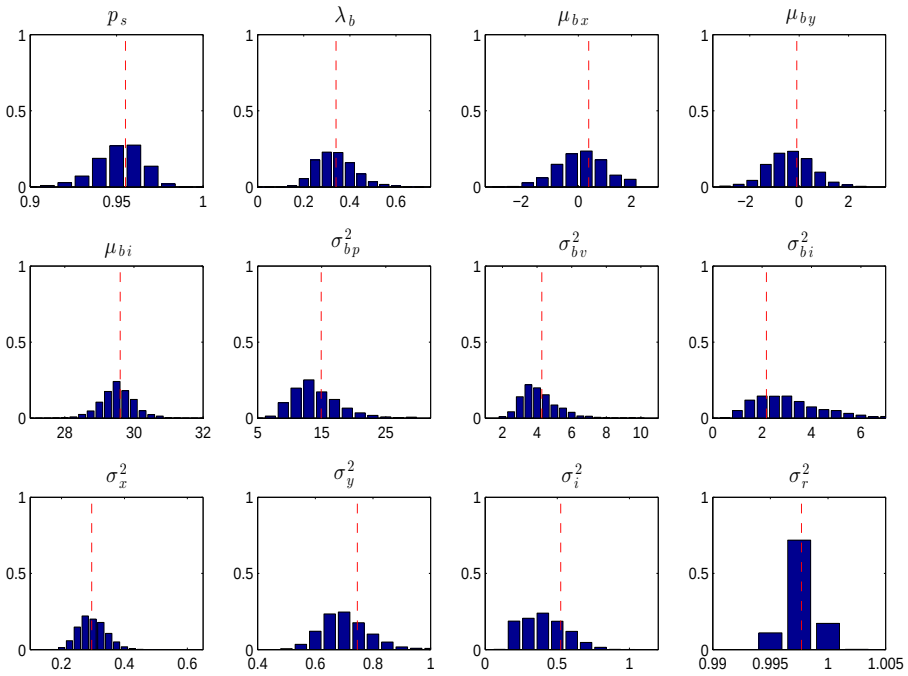

Figure 5: Histograms of 2000 parameter samples obtained by Algorithm 1. Vertical (red) lines indicate the MLE estimate in (28).

centres appear much brighter than the periphery; $\mu_{bi}^{(0)} = 70$ is calculated from (3) assuming a target is in the middle of the brightest pixel of the first frame; set $(\sigma_{bp}^2)^{(0)} = 400$ by roughly observing that the bright spots are sparse and span the whole image; $(\sigma_{bv}^2)^{(0)} = 1$ covers the velocity range of the molecules (0.1-1 pixel per image) and flat enough to allow different possible diffusion coefficients; $(\sigma_{bi}^2)^{(0)} = 100$ arbitrarily; $(\sigma_x^2)^{(0)} = (\sigma_y^2)^{(0)} = 0.1$ for small initial driving state noise; $(\sigma_i^2)^{(0)} = 25$ arbitrarily. The time-varying observation noise statistics, mean $(b_t)^{(0)}$ and variance $(\sigma_{r,t}^2)^{(0)}$, are initialised to equal the mean and variance of the pixel intensities at that time since bright pixels are sparse in the images.

The point spread parameter σ_h is not estimated and fixed at $\sigma_h = 2$ (normalised with Δ^2). The illuminated region $L(s)$ has 9×9 pixels. The value of σ_h if often known to the experimentalist otherwise, a bit tuning is required.[4] It could also be estimated as part of θ.

4.2.2 Comparison with [2]

The tracking algorithm of [2] is a nearest-neighbour method which has an image pre-processing step to extract point measurements. Tracks are then created by connecting point measurements nearest to each other. A set of consecutive frames are considered at the same time to allow temporary mis-detections. (See [2] for more details.) One of the main disadvantages of this heuristic method is that it may miss targets moving close to each other with overlapping illumination regions, as only one point measurement may be extracted from a comparably big bright region; another disadvantage is the user-defined hard-threshold which may cause the targets with lower intensities completely missed. Take frame 15 in Figure 6 as an example. Algorithm 1 detects more targets in the centre and marginal regions of the frame, compared to the method in [2]. ([2] mentions indeed targets were missed in the marginal regions.) Figure 6 compares the tracked positions of the molecules of these two algorithms. As a verification of our result, in the absence of ground truth, we compare in Figure 7 the true and synthesised image (based on our estimated tracks and model parameters) at frame $t = 15$. Their likeness offers some reassurance.

5 Conclusion

We have proposed a new MCMC based MTT algorithm for joint tracking and parameter learning that works directly with image data and avoids the need to pre-process to extract point observations. In numerical examples, we demonstrated improved performance in difficult tracking scenarios involving many targets with overlapping illumination regions, over competing methods [2, 9], which was achieved by targeting the exact posterior using MCMC. We do not advocate that our MCMC technique should replace an online method or optimised methods for non-overlapping targets such [9]. It is an alternative that works without major underlying limiting assumptions (like non-overlapping) and can be used to refine online estimates. Possible future works include the design of more efficient proposals for our MCMC routine, parellelization and performance optimization for very high density tracking.

[4] When σ_h is too small, the illuminated region taken into account is smaller than it should be which would cause many more targets than expected. In that case, we should increase σ_h.

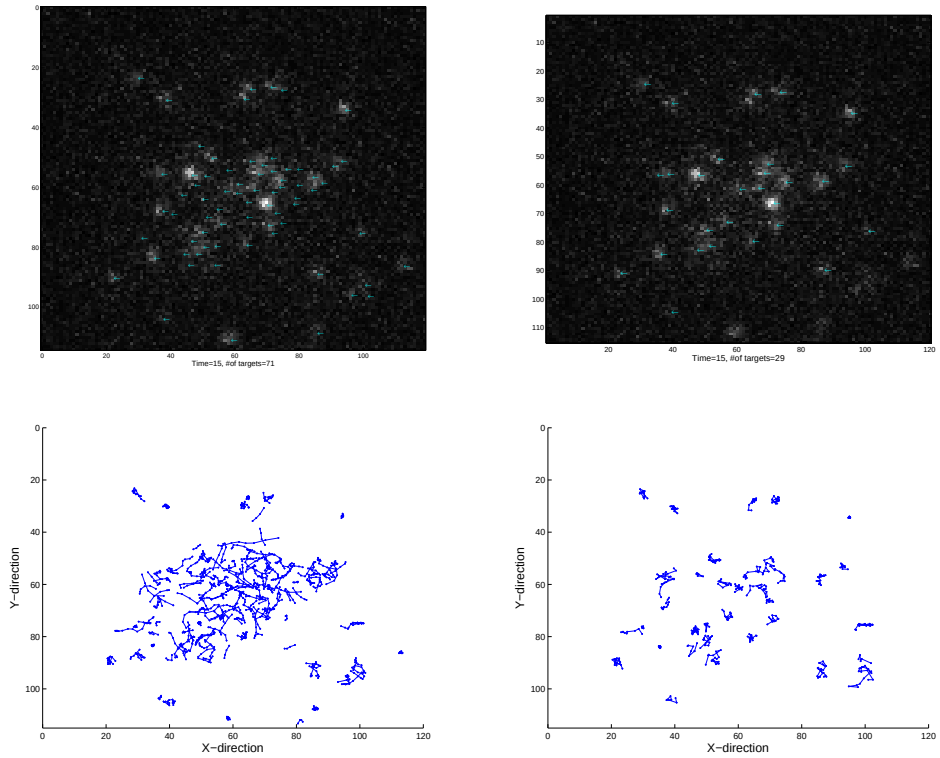

Figure 6: Estimated target positions in frame $t = 15$ of a video of fluorophore labelled molecules diffusing in the cell membrane of a Jurkat-T cell indicated by the (cyan) arrows. Top-left gives the estimated positions derived from one (representative) sample of the tracking result of Algorithm 1, top-right for the method in [2]. Bottom-left is the tracking result of our method, bottom-right the method in [2].

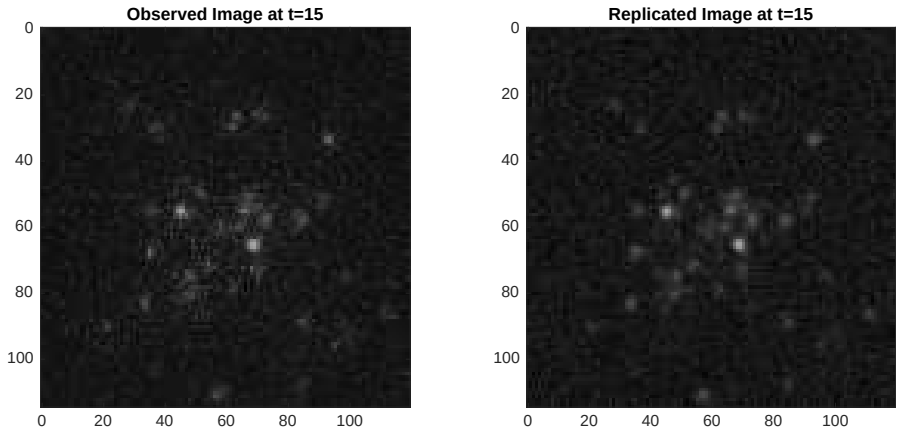

Figure 7: Comparison of the observed image (left) and the replicated one (right) at frame $t = 15$ for the real-data experiment.

A Hypothesis testing

The term of (19) to be calculated is $p(y^r_{t,L(i)}|H_1)$ where $y^r_{t,L(i)} = \{y^r_{t,j}\}_{j\in L(i)}$ and $i \in G_t$ is a pixel that is a local maximum of y^f_t. Let (r, c) be its row and column number and

$$(\bar{a}, \bar{s}) = (y^f_{t,i}, \Delta r, \Delta c).$$

Vector $(\bar{a}, \bar{s})$ can be interpreted as the likely intensity $\bar{a}$ and location $\bar{s}$ of an undetected target. Below we just write L as the set of pixels under consideration instead of $L(i)$.

Let $x = (a, s, v) \in \mathbb{R} \times \mathbb{R}^2 \times \mathbb{R}^2$ where as before a denotes intensity, $s = (s(1), s(2))$ spatial coordinates and $v = (v(1), v(2))$ spatial velocity. (Recall that a pixel illumination is not a function of velocity.) The aim is to calculate

$$\begin{aligned} p(y^r_{t,L}|H_1) &= \int \left(\prod_{j\in L} \mathcal{N}(y^r_{t,j}; a\bar{h}_j(s), \sigma^2_{r,t}) \right) \\ &\qquad \times p(a, s, v|H_1)\mathrm{d}a\mathrm{d}s\mathrm{d}v \\ &= \int p(y^r_{t,L}|a, s)p(a, s|H_1)\mathrm{d}a\mathrm{d}s \end{aligned} \tag{29}$$

where $p(a, s|H_1)$ is either the marginal (or restriction to intensity and spatial position only) of the law of the birth μ_ψ (see (1)) if proposing the initial state of the birth move, or the pdf $p(a, s|a_{0:k}, a_{0:k})$ (to be defined below) if extending the target intensity and position trajectory after having created the initial intensity and position. $p(y^r_{t,L}|a, s)$ is implicitly defined.

Let $p(y^r_{t,L}, a, s|H_1) = p(y^r_{t,L}|a, s)p(a, s|H_1)$. Use the approximation

$$\begin{aligned} \ln p(y^r_{t,L}, a, s|H_1) \approx{}& \ln p(y^r_{t,L}, \bar{a}, \bar{s}|H_1) \\ &- \frac{1}{2}[(a, s) - (\bar{a}, \bar{s})]D[(a, s) - (\bar{a}, \bar{s})]^T \end{aligned} \tag{30}$$

where $-D$ is the second order derivative $\nabla^2 \ln p(y^r_{t,L}, a, s|H_1)$ evaluated at $(\bar{a}, \bar{s})$. Expression (30) is like the Laplace approximation except that the second order Taylor expansion is computed at $(y^f_{t,i}, \Delta r, \Delta c)$ and not the true maximum $\arg\max_{a,s} p(y^r_{t,L}, a, s|H_1)$ to save on the maximization step, which we find in the numerical examples to be still effective as a component of the birth move. (Moreover, it is a fair simplification for a diffused prior.) Thus

$$p(y^r_{t,L}|H_1) \approx p(y^r_{t,L}, \bar{a}, \bar{s}|H_1)\, \frac{(2\pi)^{3/2}}{\sqrt{|D|}}$$

and the Gaussian distribution in (21) is

$$\mathcal{N}(\cdot|(\bar{a}, \bar{s}), D^{-1}). \tag{31}$$

Sec. 4 (numerical examples) assumes a linear and Gaussian model for the targets in both the synthetic and real data examples. The description is now completed by specifying $p(\bar{a}, \bar{s}|H_1)$.

When the birth move is constructing the initial/first state of the target,

$$p(a, s|H_1) = \int \mu_\psi(a, s, v)\mathrm{d}v \tag{32}$$

and the expression in (19) now simplifies to

$$\rho(y^r_{t,L(i)}) = \frac{p(H_1)}{p(H_0)} \frac{p(y^r_{t,L}|\bar{a}, \bar{s})}{p(y^r_{t,L}|H_0)} p(\bar{a}, \bar{s}|H_1) \frac{(2\pi)^{3/2}}{\sqrt{|D|}}. \tag{33}$$

Finally, we derive $p(a, s|H_1)$ in (29) when the birth move is extending the target intensity and position trajectory after having created the initial intensity position pairs $(a_0, s_0), \ldots, (a_{k-1}, s_{k-1})$ for some $k \geq 1$. For a target with a linear Gaussian model (1), the pdf of v_{k-1} (or $v_{0:k-1}$) conditioned on $s_{0:k-1}$, which is denoted $p_\psi(v_{k-1}|s_{0:k-1})$, is a Gaussian. Thus $p(a, s|H_1)$ is

$$\int f_\psi(a, s, v|a_{k-1}, s_{k-1}, v_{k-1}) p_\psi(v_{k-1}|s_{0:k-1}) \mathrm{d}v \mathrm{d}v_{k-1}$$

and the corresponding expression for $\rho(y^r_{t_k, L(i)})$ in (22) is the same as in (33).

B State proposal

The state proposal of Sec. 3.7 alters the state values of the targets whose trajectories have been partially exchanged. This proposal is defined for the Gaussian model in Sec. 4.

Let $\left\{(k, t_b^k, \mathbf{x}^k)\right\}_{k=1}^K$ be the MTT state and assume without loss of generality $U = \{1, 2\}$. The state porposal $q_{s,\theta}(U, t, \tilde{\mathbf{x}}^1, \tilde{\mathbf{x}}^2|\mathbf{x}^{1:K}, t_b^{1:K}, y_{1:n})$ can be decomposed as the product of $q_{s,\theta}(U, t|\mathbf{x}^{1:K}, t_b^{1:K})$ and $q_{s,\theta}(\tilde{\mathbf{x}}^1, \tilde{\mathbf{x}}^2|\mathbf{x}^{1:K}, t_b^{1:K}, y_{1:n}, U, t)$. The first term is the probability of selecting (U, t) which is not $y_{1:n}$ dependent. Using $y_{1:n}$ and $\left\{(k, t_b^k, \mathbf{x}^k)\right\}_{k=3}^K$, generate the residual image as in (17) by subtracting the background intensity and the contribution from all targets except $(t_b^1, \mathbf{x}^1)$ and $(t_b^2, \mathbf{x}^2)$. Let $y^r_t = (y^r_{t,1}, \ldots, y^r_{t,m})$ denote the residual images at time t. The state proposal samples $(\tilde{\mathbf{x}}^1, \tilde{\mathbf{x}}^2)$ from the pdf $q_{s,\theta}(\tilde{\mathbf{x}}^1, \tilde{\mathbf{x}}^2|\hat{\mathbf{x}}^{1:2}, \tau_b^{1:2}, y^r_{1:n})$ where $(\hat{\mathbf{x}}^1, \hat{\mathbf{x}}^2)$ is defined in (25). Note the dependancy on targets $k > 2$ is captured through the residual image.

For brevity, we write $(\mathbf{x}^1, \mathbf{x}^2)$ instead of $(\hat{\mathbf{x}}^1, \hat{\mathbf{x}}^2)$. Also $\mathbf{x}^i = (a_0^i, s_0^i, v_0^i, \ldots, a^i_{l^i-1}, s^i_{l^i-1}, v^i_{l^i-1})$ is expressed as $(\mathbf{a}^i, \mathbf{s}^i, \mathbf{v}^i)$ to highlight the intensity, spatial coordinates and velocity components.

The proposal $q_{s,\theta}$ does not alter the spatial components, i.e. $\tilde{\mathbf{s}}^1 = \mathbf{s}^1$ and $\tilde{\mathbf{s}}^2 = \mathbf{s}^2$.

The velocities $(\tilde{\mathbf{v}}^1, \tilde{\mathbf{v}}^2) \sim p_\psi(\tilde{\mathbf{v}}^1|\mathbf{s}^1) p_\psi(\tilde{\mathbf{v}}^2|\mathbf{s}^2)$, i.e. are sampled independently from p_ψ, which is the prior pdf of the velocity conditioned on the spatial coordinates values, which is Gaussian.

If targets 1 and 2 exist at time t then their state values are $x^1_{t-\tau_b^1}$ and $x^2_{t-\tau_b^2}$ and respectively. We assume (for pixel i) $y^r_{t,i} \sim \mathcal{N}(\cdot|0, \sigma^2_{r,t})$ if targets 1 and 2 both do not exist at time t, $y^r_{t,i} \sim \mathcal{N}(\cdot|a^1_{t-\tau_b^1} h_i(s^1_{t-\tau_b^1}), \sigma^2_{r,t})$ if only target 1 exists and $y^r_{t,i} \sim \mathcal{N}(\cdot|a^1_{t-\tau_b^1} h_i(s^1_{t-\tau_b^1}) + a^2_{t-\tau_b^2} h_i(s^2_{t-\tau_b^2}), \sigma^2_{r,t})$ if both exists. The prior probability model (see (1)) for the intensities are independent Gaussians, denoted $p_\psi(\mathbf{a}^1) p_\psi(\mathbf{a}^2)$. Conditioned on $y^r_{1:n}$, $(\tau_b^1, \mathbf{s}^1)$ and $(\tau_b^2, \mathbf{s}^2)$, the posterior pdf for the joint intensities is also a Gaussian, which is denoted by $q_{s,\theta}(\tilde{\mathbf{a}}^1, \tilde{\mathbf{a}}^2|\mathbf{s}^{1:2}, \tau_b^{1:2}, y^r_{1:n})$. Thus

$$\begin{aligned} q_{s,\theta}(\tilde{\mathbf{x}}^1, \tilde{\mathbf{x}}^2|\mathbf{x}^{1:2}, \tau_b^{1:2}, y^r_{1:n}) &= p_\psi(\tilde{\mathbf{v}}^1|\mathbf{s}^1) p_\psi(\tilde{\mathbf{v}}^2|\mathbf{s}^2) \\ &\quad \times q_{s,\theta}(\tilde{\mathbf{a}}^1, \tilde{\mathbf{a}}^2|\mathbf{s}^{1:2}, \tau_b^{1:2}, y^r_{1:n}). \end{aligned}$$

Acknowledgement

We thank Kristina Ganzinger and Professor David Klenerman for providing the real data and the code in [2] for the comparisons in Section 4.2, and Sinan Yıldırım for his careful reading of this paper.

References

[1] Y. Bar-Shalom and T. E. Fortmann, *Tracking and data association.* Boston: Academic Press, 1988.

[2] L. Weimann, K. A. Ganzinger, J. McColl, K. L. Irvine, S. J. Davis, N. J. Gay, C. E. Bryant, and D. Klenerman, "A quantitative comparison of single-dye tracking analysis tools using monte carlo simulations," *PloS one*, vol. 8, no. 5, p. e64287, 2013.

[3] R. P. Mahler, *Statistical multisource-multitarget information fusion*. Boston: Artech House, 2007, vol. 685.

[4] R. L. Streit, M. L. Graham, and M. J. Walsh, "Multitarget tracking of distributed targets using histogram-pmht," *Digital Signal Processing*, vol. 12, no. 2, pp. 394–404, 2002.

[5] M. G. Rutten, N. J. Gordon, and S. Maskell, "Recursive track-before-detect with target amplitude fluctuations," *IEE Proceedings - Radar, Sonar and Navigation*, vol. 152, no. 5, pp. 345–352, 2005.

[6] Y. Boers and J. Driessen, "Multitarget particle filter track before detect application," *IEE Proceedings-Radar, Sonar and Navigation*, vol. 151, no. 6, pp. 351–357, 2004.

[7] K. Punithakumar, T. Kirubarajan, and A. Sinha, "A sequential monte carlo probability hypothesis density algorithm for multitarget track-before-detect," in *Optics & Photonics 2005*. International Society for Optics and Photonics, 2005, pp. 59 131S–59 131S.

[8] S. J. Davey, M. G. Rutten, and B. Cheung, "A comparison of detection performance for several track-before-detect algorithms," *EURASIP Journal on Advances in Signal Processing*, vol. 2008, 2007.

[9] B.-N. Vo, B.-T. Vo, N.-T. Pham, and D. Suter, "Joint detection and estimation of multiple objects from image observations," *Signal Processing, IEEE Transactions on*, vol. 58, no. 10, pp. 5129–5141, 2010.

[10] F. Papi and D. Y. Kim, "A particle multi-target tracker for superpositional measurements using labeled random finite sets," *IEEE Transactions on Signal Processing*, vol. 63, no. 16, pp. 4348–4358, 2015.

[11] C. Andrieu, A. Doucet, and R. Holenstein, "Particle markov chain monte carlo methods," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 72, no. 3, pp. 269–342, 2010.

[12] L. Jiang, S. Singh, and S. Yildirim, "A new particle filtering algorithm for multiple target tracking with non-linear observations," in *Information Fusion (FUSION), 2014 17th International Conference on*, July 2014, pp. 1–8.

[13] N. Kantas, A. Doucet, S. S. Singh, J. Maciejowski, and N. Chopin, "On particle methods for parameter estimation in state-space models," *Statist. Sci.*, vol. 30, no. 3, pp. 328–351, 08 2015.

[14] S. Oh, S. Russell, and S. Sastry, "Markov chain monte carlo data association for multi-target tracking," *IEEE Trans. Automat. Control*, vol. 54, no. 3, pp. 481–497, Mar. 2009.

[15] T. Vu, B.-N. Vo, and R. Evans, "A particle marginal metropolis-hastings multi-target tracker," *IEEE Trans. Signal Process*, vol. 62, no. 15, pp. 3953 – 3964, 2014.

[16] J. Kokkala and S. Sarkka, "Combining particle MCMC with rao-blackwellized monte carlo data association for parameter estimation in multiple target tracking," *Digital Signal Processing*, vol. 47, pp. 84 – 95, 2015.

[17] L. Jiang, S. Singh, and S. Yıldırım, "Bayesian tracking and parameter learning for non-linear multiple target tracking models," *IEEE Tran. Signal Proc.*, vol. 63, pp. 5733–5745, 2015.

[18] D. Duckworth, "Monte carlo methods for multiple target tracking and parameter estimation," Master's thesis, EECS Department, University of California, Berkeley, May 2012. [Online]. Available: http://www.eecs.berkeley.edu/Pubs/TechRpts/2012/EECS-2012-68.html

[19] S. S. Singh, N. Whiteley, and S. J. Godsill, "Approximate likelihood estimation of static parameters in multi-target models," in *Bayesian Time Series Models*, D. Barber, A. T. Cemgil, and S. Chiappa, Eds. Cambridge University Press, 2011, pp. 225–244.

[20] S. Yıldırım, L. Jiang, S. S. Singh, and T. A. Dean, "Calibrating the gaussian multi-target tracking model," *Statistics and Computing*, pp. 1–14, 2014.

[21] A. Jasra, C. Holmes, and D. Stephens, "Markov chain monte carlo methods and the label switching problem in bayesian mixture modeling," *Statistical Science*, pp. 50–67, 2005.

[22] N. Whiteley, "Discussion on the paper by Andrieu, Doucet and Holenstein."

[23] P. J. Green, "Reversible jump Markov chain Monte Carlo computation and Bayesian model determination," *Biometrika*, vol. 82, no. 4, pp. 711–732, 1995.